

Adaptive Sensing Resource Allocation Over Multiple Hypothesis Tests

Dennis Wei

Abstract

This paper considers multiple binary hypothesis tests with adaptive allocation of sensing resources from a shared budget over a small number of stages. A Bayesian formulation is provided for the multistage allocation problem of minimizing the sum of Bayes risks, which is then recast as a dynamic program. In the single-stage case, the problem is a non-convex optimization, for which an algorithm composed of a series of parallel one-dimensional minimizations is presented. This algorithm ensures a global minimum under a sufficient condition. In the multistage case, the approximate dynamic programming method of open-loop feedback control is employed. In numerical simulations, the proposed allocation policies outperform alternative adaptive procedures when the numbers of true null and alternative hypotheses are not too imbalanced. In the case of few alternative hypotheses, the proposed policies are competitive using only a few stages of adaptation. In all cases substantial gains over non-adaptive sensing are observed.

Index Terms

Sequential decisions, signal detection, multiple testing, dynamic programming, non-convex optimization.

I. INTRODUCTION

This paper is concerned with the problem of multiple binary hypothesis tests under a shared sensing budget. Sensing resources can be allocated adaptively over multiple stages to the hypothesis tests, taking past observations into account. Intuitively, the advantage of adaptive allocation is that resources can be continually shifted from tests where the outcome is more certain to those that are less certain. For example, in wide-area search and surveillance, sensors can be directed to gradually concentrate more time, samples, or energy on spatial regions where target presence is the most uncertain. Other applications include adaptive spectrum sensing for unoccupied communication bands [1], biomedical clinical trials with multiple endpoints [2], and multistage gene association studies [3].

Adaptive and sequential methods for multiple hypothesis testing have been studied recently in [4]–[11]. One set of papers [4]–[6] focuses on support recovery for sparse signals. These works showed that simple multistage thresholding procedures can asymptotically drive error rates to zero with slower growth in resources compared to non-adaptive procedures; [4] focused on Gaussian observations and false discovery/non-discovery rates (FDR/FNR), while [5], [6] considered more general likelihoods and the family-wise error rate (FWER). The present work differs from and adds to [4]–[6] in three major respects: First, no sparsity assumption is made on the number of alternative (or null) hypotheses that are true. Indeed, significant performance gains are demonstrated even when the hypotheses occur in equal numbers. Second, the number of stages, i.e., the number of opportunities to adapt, is decoupled from the number of hypothesis tests and is deliberately constrained to be small. It is shown that much of the benefit of adaptation can be realized with only two or three stages. Third, a Bayesian formulation is adopted that allows for composite null and alternative hypotheses given statistical prior knowledge; [4]–[6] in contrast require a simple null hypothesis but less prior information.

A second series of works [7]–[11] has developed sequential tests that control various multiple testing error metrics: FWER [7], both type I and type II FWER simultaneously [8], [9], FDR and FNR simultaneously [10], and k -FWER or γ -FDP [11]. These procedures permit general likelihoods and dependences

between the multiple data sequences, leveraging existing methods to control sequential error rates for individual data sequences on the one hand, and the multiple testing error rates mentioned above on the other hand. Unlike [4]–[6] and this work, [7]–[11] focus on sequential procedures, which allow an indefinite number of stages at which sensing decisions can be made, subject to ensuring (perhaps conservatively) that the desired error rates are below specified levels. In contrast, in [4]–[6] and herein, both the number of stages and the resource budget are fixed while the error rates are minimized. This non-sequential setting also gives rise to the problem of resource allocation over stages and tests, which is not considered in [7]–[11].

The present paper and [4]–[11] are related more broadly to the literature on (single/non-multiple) sequential tests [12], especially with more than two hypotheses and control over observations [13]–[17]. However, while it may be possible in principle to apply these methods for more than two hypotheses to the multiple testing problem, performance losses may be expected compared to more specialized methods such as in [4]–[11] and herein. Moreover, the procedures in [13]–[17] are sequential in the sense of the previous paragraph, again in contrast to the non-sequential approach in this paper. In addition, [13]–[17] consider a finite number of sensing choices of differing quality but equal cost, whereas in this work the sensing control is continuous-valued and observation quality is a direct function of resource cost.

The statistical model and dynamic programming methods in this paper are similar to those in [18] (except for the sparsity assumption). However, the objective of hypothesis testing differs significantly from [18], which focuses on amplitude estimation of sparse signals. This difference has an important consequence for optimization: the Bayes risk adopted here as the performance metric is not a convex function of the resource allocations, unlike the estimation error metrics in [18]. The lack of convexity complicates the resource allocation problem and necessitates an alternative optimization method.

Section II presents a Bayesian formulation of multiple binary hypothesis testing with adaptive allocation of sensing resources from a fixed budget. Only Gaussian observations are considered in this paper. The multistage allocation problem of minimizing the sum of Bayes risks is then recast as a dynamic program. In Section III, single-stage and multistage solutions are developed. In the single-stage case, an algorithm is proposed involving parallel single-variable minimizations and an outer search over a Lagrange multiplier. Despite the non-convexity of the Bayes risk objective function as noted earlier, this algorithm can guarantee a global minimum when a sufficient condition is met. In the multistage case, a tractable approximate solution is proposed using open-loop feedback control [19] with the property of monotonic improvement as the number of stages increases, similar to [18]. Section IV presents numerical simulations comparing the proposed allocation policies to [4], [6], demonstrating advantages when the numbers of null and alternative hypotheses are within an order of magnitude of each other. In the highly imbalanced case, the proposed policies remain competitive and achieve most of the gains using two or three stages.

II. PROBLEM FORMULATION

We consider n binary hypothesis tests indexed by $i = 1, \dots, n$. A priori, the i th null and alternative hypotheses are true with known probabilities $\mathbb{P}(H_i = 0) = 1 - p_i(0)$ and $\mathbb{P}(H_i = 1) = p_i(0)$, and H_i, H_j are statistically independent for $i \neq j$. It is not assumed that $p_i(0) \ll 1$, i.e., the alternative hypothesis is not necessarily rare, unlike in [4]–[6].

Observations are made in T stages (indexed in parentheses) following a model similar to the one in [18]. The quality of each observation is controlled by the amount of sensing resources allocated to it. Specifically, given resource $u_i(t-1) > 0$, the observation $y_i(t)$ for test i in stage t is conditionally distributed as

$$y_i(t) \mid x_i, u_i(t-1) \sim \mathcal{N}(x_i, \nu^2/u_i(t-1)), \quad t = 1, \dots, T, \quad (1)$$

so that the precision (inverse variance) increases with $u_i(t-1)$. If $u_i(t-1) = 0$, the observation $y_i(t)$ is not taken. The mean x_i depends on H_i as specified in (3) below. The nominal variance ν^2 is assumed to be known. The observations $y_i(t)$ are independent across tests i and conditionally independent across stages t given x_i and $u_i(t)$, $t = 0, \dots, T-1$ (but not unconditionally independent).

As an example of the observation model above with $u_i(t-1)$ an integer, (1) results if $u_i(t-1)$ i.i.d. observations, each distributed as $\mathcal{N}(x_i, \nu^2)$, are taken in stage t and $y_i(t)$ is computed as the sample mean (a sufficient statistic for x_i). More generally, $u_i(t-1)$ is allowed to take on any non-negative real value to model continuous-valued resources and for mathematical convenience. The resource allocations are constrained by an overall deterministic budget,

$$\sum_{t=0}^{T-1} \sum_{i=1}^n u_i(t) = Bn, \quad (2)$$

so that the average budget per test is B . This budget constraint couples the hypothesis tests together.

In adaptive sensing, resource allocations can depend causally on all previous observations. Define $\mathbf{y}(t) = (y_1(t), \dots, y_n(t))$ (similarly for other vectors) and $\mathbf{Y}(t) = \{\mathbf{y}(1), \dots, \mathbf{y}(t)\}$. Then $u_i(t-1)$ in (1) is in general a function of $\mathbf{Y}(t-1)$. The mappings $\mathbf{Y}(t) \mapsto \mathbf{u}(t)$ are referred to as the resource allocation policy.

The mean parameters x_i in (1) are independent over i and follow Gaussian distributions conditioned on H_i ,

$$x_i \mid H_i \sim \mathcal{N}(\mu_i^{H_i}(0), \sigma_i^{H_i}(0)^2), \quad H_i = 0, 1, \quad (3)$$

with known prior parameters $\mu_i^{H_i}(0)$ and $\sigma_i^{H_i}(0)^2$. Hence both the null and alternative hypotheses can be composite if $\sigma_i^0(0), \sigma_i^1(0) > 0$, generalizing [18]. By interchanging if necessary, it is assumed that $\sigma_i^0(0) \leq \sigma_i^1(0)$ without loss of generality.

After all observations have been collected, a decision $\hat{H}_i(T) : \mathbf{Y}(T) \mapsto \{0, 1\}$ is made in each of the hypothesis tests. Performance is measured by the sum of Bayes risks,

$$R = \sum_{i=1}^n \mathbb{E}_{H_i, \mathbf{Y}(T)} \left[(1 - H_i) \hat{H}_i(T) + c H_i (1 - \hat{H}_i(T)) \right], \quad (4)$$

where $\mathbb{E}_{H_i, \mathbf{Y}(T)}$ denotes expectation over H_i and $\mathbf{Y}(T)$, and c is the cost of a Type II error (miss) relative to a Type I error (false alarm). For $c = 1$, (4) is the sum of the probabilities of error in each test, which is a union bound on the family-wise error rate, i.e., the probability of any error. It is also possible to minimize the family-wise error rate directly using an approach similar to the one herein, but this is not developed further.

In summary, the problem is to minimize the Bayes risk sum (4) with respect to the resource allocation policy $\{\mathbf{u}(t)\}$ subject to the total budget constraint (2).

A. Dynamic programming formulation

Similar to [18], the multistage minimization of the Bayes risk sum R can be cast as a dynamic program [19], where the state is a belief state summarizing the posterior distributions of H_i and x_i given observations $\mathbf{Y}(t)$. Using [18, Lem. 1] to derive these posterior distributions, it can be shown that the variables $H_i \mid \mathbf{Y}(t)$ remain independent over i with parameters $p_i(t) = \mathbb{P}(H_i = 1 \mid \mathbf{Y}(t))$, and $x_i \mid H_i, \mathbf{Y}(t)$ remain independent Gaussian with means $\mu_i^{H_i}(t) = \mathbb{E}[x_i \mid H_i, \mathbf{Y}(t)]$ and variances $\sigma_i^{H_i}(t)^2 = \text{var}(x_i \mid H_i, \mathbf{Y}(t))$. The posterior parameters evolve according to

$$p_i(t+1) = \frac{p_i(t) f_i^1(y_i(t+1); t)}{p_i(t) f_i^1(y_i(t+1); t) + (1 - p_i(t)) f_i^0(y_i(t+1); t)}, \quad (5a)$$

$$\mu_i^{H_i}(t+1) = \frac{\nu^2 \mu_i^{H_i}(t) + \sigma_i^{H_i}(t)^2 u_i(t) y_i(t+1)}{\nu^2 + \sigma_i^{H_i}(t)^2 u_i(t)}, \quad (5b)$$

$$\sigma_i^{H_i}(t+1)^2 = \frac{\nu^2 \sigma_i^{H_i}(t)^2}{\nu^2 + \sigma_i^{H_i}(t)^2 u_i(t)}, \quad (5c)$$

where in (5a), $f_i^{H_i}(\cdot; t)$ is the probability density function (PDF) of

$$y_i(t+1) \mid H_i, \mathbf{Y}(t) \sim \mathcal{N}(\mu_i^{H_i}(t), \sigma_i^{H_i}(t)^2 + \nu^2/u_i(t)). \quad (6)$$

The index $t = 0$ corresponds to the prior parameters in effect before any observations are taken.

Define the belief state as $\boldsymbol{\xi}(t) = (\mathbf{p}(t), \boldsymbol{\mu}(t), \boldsymbol{\sigma}(t)^2, U(t))$, where $\boldsymbol{\mu}(t)$ and $\boldsymbol{\sigma}(t)^2$ include all components indexed by i and $H_i = 0, 1$, and $U(t)$ is the resource budget remaining in stage t with $U(0) = Bn$. This state definition makes the objective function additive over stages, as required for a dynamic program. In fact the only explicit dependence is on the last stage, as specified below.

Proposition 1. *The Bayes risk sum (4) is the expected value of a function only of the state $\boldsymbol{\xi}(T-1)$ and control $\mathbf{u}(T-1)$,*

$$R = \sum_{i=1}^n \mathbb{E}_{\mathbf{Y}(T-1)} \left[\int_{-\infty}^{\infty} \min \{ (1 - p_i(T-1)) f_i^0(y; T-1), cp_i(T-1) f_i^1(y; T-1) \} dy \right], \quad (7)$$

where the PDFs $f_i^0(\cdot; T-1)$ and $f_i^1(\cdot; T-1)$ are completely parameterized in (6) by $\boldsymbol{\xi}(T-1)$ and $\mathbf{u}(T-1)$.

Proof: Each of the Bayes risks in (4) is minimized by the weighted maximum a posteriori (MAP) rule. Using the definition of $p_i(T)$, the i th term in (4) can thus be rewritten as

$$\mathbb{E}_{\mathbf{Y}(T)} [\min \{ 1 - p_i(T), cp_i(T) \}]. \quad (8)$$

Next we substitute for $p_i(T)$ using (5a) and iterate expectations over $y_i(T) \mid \mathbf{Y}(T-1)$ and then $\mathbf{Y}(T-1)$ to obtain

$$\mathbb{E}_{\mathbf{Y}(T-1)} \left[\int_{-\infty}^{\infty} \frac{\min \{ (1 - p_i(T-1)) f_i^0(y_i(T); T-1), cp_i(T-1) f_i^1(y_i(T); T-1) \}}{p_i(T-1) f_i^1(y_i(T); T-1) + (1 - p_i(T-1)) f_i^0(y_i(T); T-1)} \times f(y_i(T) \mid \mathbf{Y}(T-1)) dy_i(T) \right], \quad (9)$$

where the inner expectation has been expressed as an explicit integral. The denominator in (9) can be recognized as the PDF of $y_i(T) \mid \mathbf{Y}(T-1)$, yielding (7) after cancellation. ■

Remark. The allocation $\mathbf{u}(T-1)$ is also constrained by the remaining budget $U(T-1)$, which is part of the state $\boldsymbol{\xi}(T-1)$. An equivalent unconstrained formulation can be obtained by augmenting (7) with the stipulation that R is infinite if $\sum_{i=1}^n u_i(T-1) > U(T-1)$, i.e., the budget is exceeded.

III. RESOURCE ALLOCATION POLICIES

This section discusses single-stage and multistage resource allocation policies that minimize the Bayes risk sum (7) under the budget constraint (2). As discussed in Section III-B, the single-stage policy of Section III-A also applies to the last stage of any multistage policy.

A. Single-stage policy

In the single-stage case $T = 1$, the expectation in (7) is absent and the objective function simplifies. The remaining integral is the Bayes risk of the optimal test between two Gaussian distributions with different means and variances. Let $R_i(u_i; \boldsymbol{\xi}_i)$ denote this Bayes risk, where the stage index $T-1$ is suppressed to simplify notation, and $\boldsymbol{\xi}_i$ represents the components of the state with index i . Appendix A provides explicit expressions for $R_i(u_i; \boldsymbol{\xi}_i)$ in terms of the standard Gaussian cumulative distribution function (CDF). The single-stage resource allocation problem is therefore

$$R^*(\xi) = \min_{\mathbf{u}} \sum_{i=1}^n R_i(u_i; \xi_i) \quad \text{s.t.} \quad \sum_{i=1}^n u_i = U, \quad u_i \geq 0 \quad \forall i. \quad (10)$$

Fig. 1(a) shows that the Bayes risk $R_i(u_i; \xi_i)$ is a decreasing but non-convex function of u_i for a particular choice of parameters ξ_i . These properties hold in general for other choices of ξ_i , implying that (10) is a non-convex optimization problem.




Fig. 1. (a) The Bayes risk $R_i(u_i; \xi_i)$ is a decreasing but non-convex function of u_i . (b) The Lagrangian in (11) can have more than one minimizer.

Despite the absence of convexity, it is still possible in some cases to guarantee a globally optimal solution to (10). We consider minimizing a Lagrangian of (10) in which only the equality constraint is dualized with Lagrange multiplier λ . The Lagrangian then decouples over i . Define the (possibly non-unique) minimizer of each Lagrangian component as

$$u_i(\lambda) \in \arg \min_{u_i \geq 0} R_i(u_i; \xi_i) + \lambda u_i. \quad (11)$$

Since $R_i(u_i; \xi_i)$ is bounded from above by $\min\{1 - p_i, cp_i\} = R_i(0; \xi_i)$, a negative value for λ in (11) would result in divergence toward infinity. Hence it is sufficient to consider $\lambda \geq 0$. The following result gives a sufficient condition for $\mathbf{u}(\lambda) = (u_1(\lambda), \dots, u_n(\lambda))$ to be globally optimal for (10).

Proposition 2. *If there exists a Lagrange multiplier $\lambda \geq 0$ such that a set of minimizers $\mathbf{u}(\lambda) = (u_1(\lambda), \dots, u_n(\lambda))$ defined by (11) is feasible for problem (10), then $\mathbf{u}(\lambda)$ is a global minimum of (10).*

Proof: This is an adaptation of [20, Prop. 3.3.4], where the equality constraint $\sum_{i=1}^n u_i = U$ in (10) (or equivalently two inequality constraints) has been incorporated into the Lagrangian function, while the remaining constraint set X is the non-negative orthant. ■

The minimization in (11) also satisfies the monotonicity property below, which confirms the interpretation of λ as a penalty parameter.

Lemma 1. *If $\lambda_1 < \lambda_2$, then $u_i(\lambda_1) \geq u_i(\lambda_2)$ for any minimizers $u_i(\lambda_1), u_i(\lambda_2)$ in (11).*

Proof: Let $u_i(\lambda_1)$ be any minimizer in (11) for $\lambda = \lambda_1$. Then for all $u_i > u_i(\lambda_1)$,

$$R_i(u_i(\lambda_1); \xi_i) + \lambda_1 u_i(\lambda_1) \leq R_i(u_i; \xi_i) + \lambda_1 u_i. \quad (12)$$

By assumption, we have

$$(\lambda_2 - \lambda_1)u_i(\lambda_1) < (\lambda_2 - \lambda_1)u_i. \quad (13)$$

Adding (12) and (13) yields

$$R_i(u_i(\lambda_1); \xi_i) + \lambda_2 u_i(\lambda_1) < R_i(u_i; \xi_i) + \lambda_2 u_i,$$

which implies that any minimizer $u_i(\lambda_2)$ of (11) for $\lambda = \lambda_2$ must be no greater than $u_i(\lambda_1)$. ■

Based on Proposition 2 and Lemma 1, the following algorithm is proposed to solve (10), consisting of an outer bisection search over λ and inner single-variable minimizations (11) to determine $u_i(\lambda)$, $i = 1, \dots, n$, which can be done in parallel. Lower and upper bounds \underline{u}_i and \bar{u}_i are maintained on each u_i , where initially $\underline{u}_i = 0$ and $\bar{u}_i = \infty$. Any algorithm can be used to solve (11) subject to the bounds $\underline{u}_i \leq u_i \leq \bar{u}_i$, for example gradient descent with logarithmically-spaced line search as used to generate the results in Section IV. Let $S(\lambda) = \sum_{i=1}^n u_i(\lambda)$. If for a given λ , the resulting $u_i(\lambda)$ satisfy $S(\lambda) < U$, then λ is decreased according to the bisection method, the lower bounds \underline{u}_i are updated to the current solutions $u_i(\lambda)$, exploiting Lemma 1, and (11) is re-solved. Analogous actions are taken if $S(\lambda) > U$. If $S(\lambda) = U$, then by Proposition 2, the algorithm terminates with a globally optimal solution to (10).

For the bisection search over λ , the initial lower bound is set at 0. The lemma below is used to set the initial upper bound.

Lemma 2. *Any minimizer $u_i(\lambda)$ in (11) is bounded from above as $u_i(\lambda) < R_i(0; \xi_i)/\lambda = \min\{1 - p_i, cp_i\}/\lambda$.*

Proof: Since the Bayes risk $R_i(u_i; \xi_i)$ is positive for finite u_i , if $u_i \geq R_i(0; \xi_i)/\lambda$ then $R_i(u_i; \xi_i) + \lambda u_i > R_i(0; \xi_i)$ and u_i cannot be minimal. ■

It follows that a sufficient upper bound on λ is $\sum_{i=1}^n R_i(0; \xi_i)/U$, since any higher value can be seen to result in $S(\lambda) < U$. Lemma 2 is also used to further constrain the inner minimizations over u_i when it gives a tighter upper bound than \bar{u}_i .

The above algorithm does not always ensure a global minimum for (10). Specifically, it may not be possible to satisfy the condition in Proposition 2, i.e., there is no λ for which $S(\lambda) = U$ to make $\mathbf{u}(\lambda)$ feasible. The problem is illustrated in Fig. 1(b), which shows a value for λ such that the Lagrangian in (11) has two separated minimizers. Any change in λ would result in either the left or the right minimizer being unique. Hence the function $S(\lambda)$ is discontinuous and the bisection search over λ may not converge with $S(\lambda) = U$. For the numerical results in Section IV, cases of non-convergence are addressed simply by rescaling the final solution $\mathbf{u}(\lambda)$ so that it sums to U . The loss in optimality appears to be insignificant for large n and can even be bounded analytically, although this is not presented here.

B. Multistage policies

In a multistage adaptive policy, the last-stage allocation $\mathbf{u}(T-1)$ can depend on all previous observations $\mathbf{Y}(T-1)$. In other words, $\mathbf{u}(T-1)$ is determined after conditioning on $\mathbf{Y}(T-1)$, which again removes the expectation from (7). Therefore the last-stage allocation problem in any multistage policy reduces to the single-stage case (10).

For a two-stage policy, it remains to determine the first-stage allocation $\mathbf{u}(0)$. This is done recursively by solving

$$\min_{\mathbf{u}(0)} \mathbb{E}_{\mathbf{y}(1)} [R^*(\xi(1)) \mid \xi(0), \mathbf{u}(0)] \quad \text{s.t.} \quad \sum_{i=1}^n u_i(0) \leq U(0), \quad u_i(0) \geq 0 \quad \forall i, \quad (14)$$

where $R^*(\xi(1))$ is defined by (10) as the optimal cost of the second stage, and the conditional notation reflects the parameterization of the distribution of $\mathbf{y}(1)$ in terms of $\xi(0)$ and $\mathbf{u}(0)$ (see (6)). In the case of priors that are homogeneous over i , i.e., $p_i(0)$, $\mu_i^{H_i}(0)$, $\sigma_i^{H_i}(0)^2$ do not depend on i (but can depend on H_i), then the first-stage allocation is also homogeneous by symmetry, $u_i(0) = u(0)$, and (14) becomes a scalar minimization with respect to $u(0) \in [0, U(0)/n]$. This minimization is performed offline using Monte Carlo samples of $\mathbf{y}(1)$ to approximate the expectation in (14) and the algorithm in Section III-A to approximate $R^*(\xi(1))$ for each realization of $\mathbf{y}(1)$.

For an inhomogeneous prior or more than two stages, an open-loop feedback control (OLFC) policy [19] is employed, similar to [18]. Consider the problem of determining the allocation $\mathbf{u}(t)$ in stage $t < T-1$ conditioned on available observations $\mathbf{Y}(t)$ through the state $\xi(t)$. In exact dynamic programming, $\mathbf{u}(t)$ is

optimized assuming that future allocations $\mathbf{u}(t+1), \dots, \mathbf{u}(T-1)$ are also chosen optimally as functions of $\mathbf{Y}(t+1), \dots, \mathbf{Y}(T-1)$ respectively. However, in stage t these future observations are not available and are therefore random quantities, which greatly complicates the optimization. The OLFC simplification is to assume that $\mathbf{u}(t+1), \dots, \mathbf{u}(T-1)$ can depend only on current observations $\mathbf{Y}(t)$, i.e., future planning is done “open-loop”. This leads to a joint optimization over $\mathbf{u}(t), \mathbf{u}(t+1), \dots, \mathbf{u}(T-1)$ of the Bayes risk sum (4) conditioned on $\boldsymbol{\xi}(t)$:

$$\min_{\mathbf{u}(t), \dots, \mathbf{u}(T-1)} \sum_{i=1}^n \mathbb{E} [\min\{1 - p_i(T), cp_i(T)\} \mid \boldsymbol{\xi}(t)] \quad \text{s.t.} \quad \sum_{i=1}^n \sum_{\tau=t}^{T-1} u_i(\tau) = U(t), \quad u_i(\tau) \geq 0 \quad \forall i, \tau, \quad (15)$$

where (8) has been substituted into the objective function. Once (15) is solved, only the first stage $\mathbf{u}(t)$ is applied to collect new observations $\mathbf{y}(t+1)$ as in (1) and update the state to $\boldsymbol{\xi}(t+1)$ using (5). Then (15) is solved for $\mathbf{u}(t+1), \dots, \mathbf{u}(T-1)$ given $\boldsymbol{\xi}(t+1)$ under the same OLFC assumption, and the process continues.

The OLFC optimization problem (15) can be further simplified to an instance of the single-stage optimization (10). This together with Appendix A provides an explicit expression for the objective function in terms of Gaussian CDFs, i.e. without expectation operators, and also reduces the number of optimization variables from $n(T-t)$ to n .

Lemma 3. *Let*

$$v_i(t) = \sum_{\tau=t}^{T-1} u_i(\tau).$$

The OLFC optimization problem (15) reduces to an instance of the single-stage optimization (10) with $u_i = v_i(t)$, $\boldsymbol{\xi} = \boldsymbol{\xi}(t)$, and $U = U(t)$.

Proof: The first step is to relate the Bayes risk objective in (15) to the state $\boldsymbol{\xi}(t)$ in stage t . Recalling the definition $p_i(t) = \mathbb{P}(H_i = 1 \mid \mathbf{Y}(t))$, an application of Bayes rule similar to (5a) yields

$$p_i(t) = \frac{p_i(t)f(y_i(t+1), \dots, y_i(T) \mid H_i = 1, \mathbf{Y}(t))}{f(y_i(t+1), \dots, y_i(T) \mid \mathbf{Y}(t))},$$

$$1 - p_i(t) = \frac{(1 - p_i(t))f(y_i(t+1), \dots, y_i(T) \mid H_i = 0, \mathbf{Y}(t))}{f(y_i(t+1), \dots, y_i(T) \mid \mathbf{Y}(t))}.$$

Hence

$$\begin{aligned} \mathbb{E} [\min\{1 - p_i(T), cp_i(T)\} \mid \boldsymbol{\xi}(t)] &= \int \dots \int \min \left\{ \frac{(1 - p_i(t))f(y_i(t+1), \dots, y_i(T) \mid H_i = 0, \boldsymbol{\xi}(t))}{f(y_i(t+1), \dots, y_i(T) \mid \boldsymbol{\xi}(t))}, \right. \\ &\quad \left. \frac{cp_i(t)f(y_i(t+1), \dots, y_i(T) \mid H_i = 1, \boldsymbol{\xi}(t))}{f(y_i(t+1), \dots, y_i(T) \mid \boldsymbol{\xi}(t))} \right\} \\ &\quad \times f(y_i(t+1), \dots, y_i(T) \mid \boldsymbol{\xi}(t)) dy_i(t+1) \dots dy_i(T) \\ &= \int \dots \int \min \{ (1 - p_i(t))f(y_i(t+1), \dots, y_i(T) \mid H_i = 0, \boldsymbol{\xi}(t)), \\ &\quad cp_i(t)f(y_i(t+1), \dots, y_i(T) \mid H_i = 1, \boldsymbol{\xi}(t)) \} dy_i(t+1) \dots dy_i(T). \end{aligned} \quad (16)$$

To simplify (16), a Neyman factorization is derived for the joint density $f(y_i(t+1), \dots, y_i(T) \mid H_i, \boldsymbol{\xi}(t))$. Toward this end, we have

$$f(y_i(t+1), \dots, y_i(T) \mid H_i, \boldsymbol{\xi}(t)) = \int f(y_i(t+1), \dots, y_i(T) \mid x_i, \boldsymbol{\xi}(t)) f(x_i \mid H_i, \boldsymbol{\xi}(t)) dx_i. \quad (17)$$

Under the OLFC assumption, conditioning on $\xi(t)$ also fixes the allocations $u_i(t), \dots, u_i(T-1)$. Therefore $y_i(t+1), \dots, y_i(T) \mid x_i, \xi(t)$ are independent Gaussian according to (1). Furthermore, it is straightforward to show that the weighted average

$$\bar{y}_i(t) \equiv \frac{\sum_{\tau=t}^{T-1} u_i(\tau) y_i(\tau+1)}{\sum_{\tau=t}^{T-1} u_i(\tau)}$$

is distributed as

$$\bar{y}_i(t) \mid x_i, \xi(t) \sim \mathcal{N}\left(x_i, \frac{\nu^2}{\sum_{\tau=t}^{T-1} u_i(\tau)}\right) = \mathcal{N}\left(x_i, \frac{\nu^2}{v_i(t)}\right)$$

and is a sufficient statistic for x_i . It follows that (17) can be rewritten as

$$\begin{aligned} f(y_i(t+1), \dots, y_i(T) \mid H_i, \xi(t)) &= f(y_i(t+1), \dots, y_i(T) \mid \bar{y}_i(t), \xi(t)) \\ &\quad \times \int f(\bar{y}_i(t) \mid x_i, \xi(t)) f(x_i \mid H_i, \xi(t)) dx_i \\ &= f(y_i(t+1), \dots, y_i(T) \mid \bar{y}_i(t), \xi(t)) f(\bar{y}_i(t) \mid H_i, \xi(t)), \end{aligned} \quad (18)$$

where

$$\bar{y}_i(t) \mid H_i, \xi(t) \sim \mathcal{N}(\mu_i^{H_i}(t), \sigma_i^{H_i}(t)^2 + \nu^2/v_i(t)) \quad (19)$$

as a result of compounding, similar to (6).

The final step is to substitute the factorization (18) into (16). Upon doing so, it is seen that the common factor $f(y_i(t+1), \dots, y_i(T) \mid \bar{y}_i(t), \xi(t))$ integrates to 1, leaving

$$\begin{aligned} \mathbb{E}[\min\{1 - p_i(T), cp_i(T)\} \mid \xi(t)] &= \int \min\{(1 - p_i(t)) f(\bar{y}_i(t) \mid H_i = 0, \xi(t)), \\ &\quad cp_i(t) f(\bar{y}_i(t) \mid H_i = 1, \xi(t))\} d\bar{y}_i(t). \end{aligned}$$

Comparing the above expression with $R_i(u_i; \xi_i)$, defined as the integral in (7), and (19) with (6), we conclude that

$$\mathbb{E}[\min\{1 - p_i(T), cp_i(T)\} \mid \xi(t)] = R_i(v_i(t); \xi_i(t)).$$

Rewriting the constraint in (15) in terms of $v_i(t)$ completes the proof. \blacksquare

According to Lemma 3, the OLFC allocation in stage t can be determined by first solving the single-stage problem (10) with appropriate parameters. However, the resulting solution $\mathbf{v}^*(t)$ does not specify the allocations of the sums $v_i(t) = \sum_{\tau=t}^{T-1} u_i(\tau)$ over stages, in particular the first allocation $\mathbf{u}(t)$ used to make new observations. For this purpose, the approach in [18] is followed in which $\mathbf{u}(t)$ is constrained to be a scaled version of $\mathbf{v}^*(t)$: $\mathbf{u}(t) = \beta(t; T) \mathbf{v}^*(t)$ where $\beta(t; T) \in [0, 1]$ and the second argument T denotes the number of stages in the policy. Setting $\beta(t; T) < 1$ thus conserves some of the resource budget for future stages.

The multipliers $\beta(t; T)$ are determined recursively for $T = 1, 2, \dots$ as follows. For $t = T-1$, $\mathbf{v}^*(T-1)$ coincides with $\mathbf{u}(T-1)$ and $\beta(T-1; T) = 1$. This case encompasses the single-stage ($T = 1$) policy described in Section III-A and the last-stage allocation discussed at the beginning of Section III-B. For $T > 1$, multipliers are reused across policies with different numbers of stages to reduce the number of degrees of freedom. Specifically,

$$\beta(t; T) = \beta(t-1; T-1), \quad t = 1, 2, \dots, T-2. \quad (20)$$

The remaining first-stage multiplier $\beta(0; T)$ is optimized in a manner similar to (14). Define $R^{\text{OLFC}-T}(\xi(1))$ to be the Bayes risk cost of a T -stage OLFC allocation policy starting from stage 1 and belief state $\xi(1)$ and using multipliers equal to those of a previously determined $T-1$ -stage policy (20). Then

$$\beta(0; T) = \arg \min_{\beta \in [0, 1]} \mathbb{E}[R^{\text{OLFC}-T}(\xi(1)) \mid \xi(0), \beta \mathbf{v}^*(0)]. \quad (21)$$

This one-dimensional optimization can be carried out offline using Monte Carlo samples both to approximate the expectation as well as to simulate the cost $R^{\text{OLFC}}(\xi(1))$ of the policy from stage 1 onward. As shown in [18, Prop. 2], an important property of the procedure summarized by (20)–(21) is that the resulting OLFC policies improve monotonically with the number of stages T . Further details can be found in [18].

IV. NUMERICAL RESULTS

The multistage resource allocation policies described in Section III are numerically compared to the distilled sensing (DS) [4] and sequential thresholding (ST) [6] procedures, as well as to a single-stage non-adaptive baseline policy (NA). For the results presented below, the number of hypothesis tests n is 10^4 and a homogeneous prior is used: $p_i(0) = p(0)$, $\mu_i^0(0) = 0$, $\mu_i^1(0) = 1$, $\sigma_i^0(0)^2 = 0$, and $\sigma_i^1(0)^2 = 1/16$ for all i . Observations are simulated according to (1) and (3). The observation noise parameter ν^2 is normalized to 1 and the average budget per test B is varied. Since ν^2 and $u_i(t)$ always appear in the same ratio as in (1), an equivalent alternative would be to fix B and vary ν^2 instead. The performance metric is (4) with $c = 1$, i.e., it is the expected number of errors of either type.

The number of stages in the proposed OLFC policies is limited between 2 and 4. In all cases, the first-stage allocation $u(0)$ is uniform because of the homogeneous prior. For $T = 2$, Fig. 2(a) shows the first-stage budget fraction $u(0)$ that results from the offline optimization (14) for different values of $p(0)$ and B . Performance is not too sensitive to the exact value of $u(0)$ since the objective function in (14) tends to be relatively flat away from the extremes $u(0) = 0$ and $u(0) = 1$. Fig. 2(b) plots the same parameter $u(0)$ for $T = 3$.




Fig. 2. First-stage allocation $u_i(0) = u(0)$ in the proposed 2-stage (a) and 3-stage (b) policies as a function of the mean proportion $p(0)$ of alternative hypotheses and the resource budget per test B .

For DS and ST, while [4], [6] prescribe values for T as functions of n , in these experiments all $T \in \{2, \dots, 12\}$ are tested and results for the best T are shown. A similar optimization is performed over the parameter $\rho \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ in [6]. The budget allocations over stages follow [4, eq. (4),(5)] and [6, eq. (14)] respectively, except in the last stage of ST where the remaining budget is used up entirely. Two versions of DS and ST are implemented: the versions originally proposed in [4], [6] that use only the last stage of observations to make decisions, and Bayesian versions (DSB, STB), not proposed in [4], [6], in which the allocations $u(t)$ are specified by [4], [6] but inference is done through the posterior update equations (5), thus incorporating all stages of observations. As seen below, the Bayesian versions perform considerably better.

The performance of the policies is compared in Fig. 3. For equiprobable hypotheses, $p(0) = 0.5$, the proposed 2-stage policy achieves significant reductions in error (up to a factor of 5) relative to the

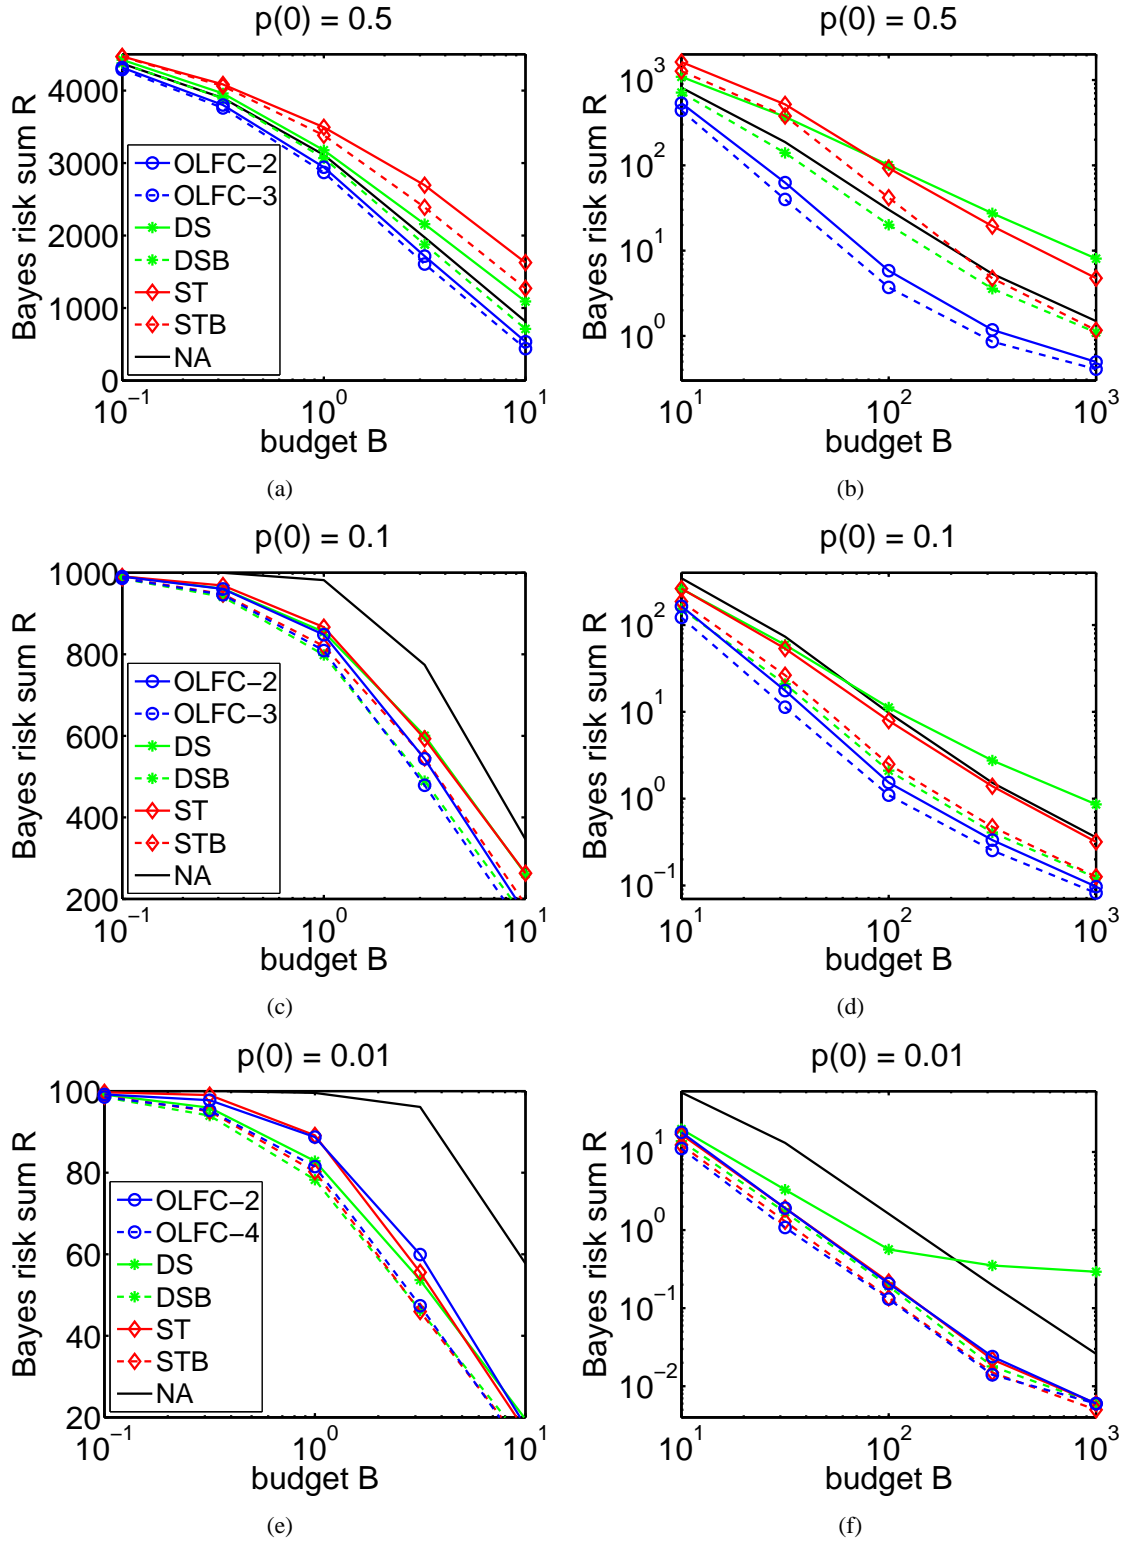


Fig. 3. Expected number of errors ((4) with $c = 1$) resulting from the proposed open-loop feedback control policies with T stages (OLFC- T), original and Bayesian versions of distilled sensing (DS, DSB) and sequential thresholding (ST, STB), and a non-adaptive baseline (NA). The legends in (a), (c), (e) also apply to (b), (d), (f) respectively. For $p(0) = 0.5, 0.1$ in (a)–(d), OLFC-2 and/or OLFC-3 outperform the alternative methods across budget levels. For $p(0) = 0.01$ in (e)(f), OLFC-4 is competitive with DSB and STB, while OLFC-2 achieves most of the gains using only 2 stages.

baseline NA policy, while the 3-stage OLFC policy yields further improvement. Since DS(B) and ST(B) are not designed for this non-sparse scenario, they perform less well, in some cases worse than NA. For

$p(0) = 0.1$, the 3-stage OLFC policy essentially dominates the other policies, and at moderate to large resource levels in Fig. 3(d), it is joined by the 2-stage OLFC policy. For $p(0) = 0.01$ and low resources in Fig. 3(e), DSB and STB have slightly lower error rates than the 4-stage OLFC policy, while for higher resources in Fig. 3(f), the opposite is true. Moreover, the 2-stage OLFC policy attains most of the gains of these best-performing policies that use more stages. In particular, the optimized DSB and STB policies shown in Fig. 3 for $B \leq 1$ use at least 8 and 6 stages respectively.

V. CONCLUSION

This paper has explored the benefits of adaptive sensing for multiple binary hypothesis testing, notably in the regimes of balanced null and alternative hypotheses and few allocation stages. Future work includes generalizations to non-Gaussian observations, refinements of both the single-stage optimization and multistage dynamic programming procedures, and theoretical analysis that aims especially to understand the gains in the non-sparse setting and at moderate, non-asymptotic resource levels.

APPENDIX A BAYES RISK COMPUTATION

This appendix derives the optimal Bayes risk for two Gaussian distributions, i.e., the integral in (7) denoted as $R_i(u_i; \xi_i)$ in Section III-A.

To simplify notation in this appendix, both the stage index $T - 1$ and test index i are dropped. Furthermore, we define $\mu \equiv \mu^1 - \mu^0$ and shift the distributions so that $\mu^0 = 0$ and $\mu^1 = \mu$, without changing the Bayes risk. Define $\Sigma^{H_i} = (\sigma^{H_i})^2 + \nu^2/u$, $H_i = 0, 1$, to be the conditional variances in (6). Recalling from Section II the assumption that $\sigma_i^0(0) \leq \sigma_i^1(0)$, it can be seen that $\Sigma^0 \leq \Sigma^1$. Two cases are considered.

Case $\Sigma^0 < \Sigma^1$: First the decision regions

$$\begin{aligned}\mathcal{Y}^0 &= \{(1-p)f^0(y) \geq cpf^1(y)\}, \\ \mathcal{Y}^1 &= \{(1-p)f^0(y) < cpf^1(y)\}\end{aligned}$$

are determined, corresponding to the two terms in the minimization in (7). Taking logarithms and collecting terms gives the following quadratic inequalities for \mathcal{Y}_0 :

$$\begin{aligned}\log(1-p) - \frac{1}{2} \log \Sigma^0 - \frac{y^2}{2\Sigma^0} &\geq \log(cp) - \frac{1}{2} \log \Sigma^1 - \frac{(y-\mu)^2}{2\Sigma^1}, \\ \left(\frac{1}{2\Sigma^0} - \frac{1}{2\Sigma^1}\right) y^2 + \frac{\mu}{\Sigma^1} y - \frac{\mu^2}{2\Sigma^1} - \frac{1}{2} \log \frac{\Sigma^1}{\Sigma^0} - \log \frac{1-p}{cp} &\leq 0,\end{aligned}\tag{A.1}$$

with the inequalities reversed for \mathcal{Y}_1 . Applying the quadratic formula to (A.1) yields the decision boundaries

$$\begin{aligned}y_{\pm} &= \frac{\Sigma^0 \Sigma^1}{\Sigma^1 - \Sigma^0} \left(-\frac{\mu}{\Sigma^1} \pm \sqrt{\left(\frac{\mu}{\Sigma^1}\right)^2 + \frac{\Sigma^1 - \Sigma^0}{\Sigma^0 \Sigma^1} \left(\frac{\mu^2}{\Sigma^1} + \log \frac{\Sigma^1}{\Sigma^0} + 2 \log \frac{1-p}{cp}\right)} \right) \\ &= \frac{-\Sigma^0 \mu \pm \sqrt{\Sigma^0 \Sigma^1 D}}{\Sigma^1 - \Sigma^0},\end{aligned}$$

provided that the discriminant

$$D = \mu^2 + (\Sigma^1 - \Sigma^0) \left(\log \frac{\Sigma^1}{\Sigma^0} + 2 \log \frac{1-p}{cp} \right)$$

is non-negative. The region \mathcal{Y}^0 is the interval $[y_-, y_+]$ while the region \mathcal{Y}^1 is the union of intervals $(-\infty, y_-) \cup (y_+, \infty)$. If $D < 0$, then the decision boundaries do not exist, $\mathcal{Y}^0 = \emptyset$, and $\mathcal{Y}^1 = \mathbb{R}$.

Next the integrals of $f^0(y)$ and $f^1(y)$ are evaluated over \mathcal{Y}^1 and \mathcal{Y}^0 respectively, corresponding to the Type I and Type II error probabilities. By standardizing the decision boundaries y_{\pm} , the Type I error probability can be expressed in terms of the standard Gaussian CDF Φ as

$$\mathbb{P}^0(\mathcal{Y}^1) = \Phi\left(-\frac{y_+}{\sqrt{\Sigma^0}}\right) + \Phi\left(\frac{y_-}{\sqrt{\Sigma^0}}\right) = \Phi\left(\frac{\sqrt{\Sigma^0}\mu - \sqrt{\Sigma^1 D}}{\Sigma^1 - \Sigma^0}\right) + \Phi\left(\frac{-\sqrt{\Sigma^0}\mu - \sqrt{\Sigma^1 D}}{\Sigma^1 - \Sigma^0}\right).$$

Similarly the Type II error probability is

$$\mathbb{P}^1(\mathcal{Y}^0) = \Phi\left(\frac{y_+ - \mu}{\sqrt{\Sigma^1}}\right) - \Phi\left(\frac{y_- - \mu}{\sqrt{\Sigma^1}}\right) = \Phi\left(\frac{-\sqrt{\Sigma^1}\mu + \sqrt{\Sigma^0 D}}{\Sigma^1 - \Sigma^0}\right) - \Phi\left(\frac{-\sqrt{\Sigma^1}\mu - \sqrt{\Sigma^0 D}}{\Sigma^1 - \Sigma^0}\right).$$

The Bayes risk is then given by the linear combination

$$R_i(u; \xi) = (1 - p) \mathbb{P}^0(\mathcal{Y}^1) + cp \mathbb{P}^1(\mathcal{Y}^0). \quad (\text{A.2})$$

Case $\Sigma^0 = \Sigma^1 = \Sigma$: In this case (A.1) simplifies to

$$y \leq \frac{\mu}{2} + \frac{\Sigma}{\mu} \log \frac{1-p}{cp} \equiv y_c$$

for region \mathcal{Y}^0 , and $y > y_c$ for region \mathcal{Y}^1 . The error probabilities are therefore

$$\begin{aligned} \mathbb{P}^0(\mathcal{Y}^1) &= \Phi\left(-\frac{y_c}{\sqrt{\Sigma}}\right) = \Phi\left(-\frac{\mu}{2\sqrt{\Sigma}} - \frac{\sqrt{\Sigma}}{\mu} \log \frac{1-p}{cp}\right), \\ \mathbb{P}^1(\mathcal{Y}^0) &= \Phi\left(\frac{y_c - \mu}{\sqrt{\Sigma}}\right) = \Phi\left(-\frac{\mu}{2\sqrt{\Sigma}} + \frac{\sqrt{\Sigma}}{\mu} \log \frac{1-p}{cp}\right), \end{aligned}$$

and the Bayes risk is still given by (A.2).

REFERENCES

- [1] A. Tajer, R. M. Castro, and X. Wang, "Adaptive sensing of congested spectrum bands," *IEEE Trans. Inf. Theory*, vol. 58, no. 9, pp. 6110–6125, Sep. 2012.
- [2] C. Jennison and B. W. Turnbull, *Group Sequential Methods with Applications to Clinical Trials*. New York: Chapman & Hall/CRC, 2000.
- [3] S. Zehetmayer, P. Bauer, and M. Posch, "Optimized multi-stage designs controlling the false discovery or the family-wise error rate," *Stat. Med.*, vol. 27, pp. 4145–4160, 2008.
- [4] J. Haupt, R. M. Castro, and R. Nowak, "Distilled sensing: Adaptive sampling for sparse detection and estimation," *IEEE Trans. Inf. Theory*, vol. 57, pp. 6222–6235, Sep. 2011.
- [5] M. Malloy and R. Nowak, "On the limits of sequential testing in high dimensions," in *Conf. Rec. Asilomar Conf. Signals Syst. Comput.*, Nov. 2011, pp. 1245–1249.
- [6] M. L. Malloy and R. D. Nowak, "Sequential testing for sparse recovery," Dec. 2012, arXiv:1212.1801v1.
- [7] J. Bartroff and T. L. Lai, "Multistage tests of multiple hypotheses," *Commun. Stat. Theory Methods*, vol. 39, no. 8-9, pp. 1597–1607, 2010.
- [8] S. K. De and M. Baron, "Step-up and step-down methods for testing multiple hypotheses in sequential experiments," *J. Stat. Plan. Inference*, vol. 142, no. 7, pp. 2059–2070, Jul. 2012.
- [9] J. Bartroff and J. Song, "Sequential tests of multiple hypotheses controlling type I and II familywise error rates," *J. Stat. Plan. Inference*, vol. 153, pp. 100–114, Oct. 2014.
- [10] —, "Sequential tests of multiple hypotheses controlling false discovery and nondiscovery rates," Nov. 2013, arXiv:1311.3350.
- [11] J. Bartroff, "Multiple hypothesis tests controlling generalized error rates for sequential data," Sep. 2014, arXiv:1406.5933.
- [12] A. Wald and J. Wolfowitz, "Optimum character of the sequential probability ratio test," *Ann. Math. Stat.*, vol. 19, no. 3, pp. 326–339, 1948.
- [13] H. Chernoff, "Sequential design of experiments," *Ann. Math. Stat.*, vol. 30, pp. 755–770, 1959.
- [14] S. A. Bessler, "Theory and applications of the sequential design of experiments, k-actions and infinitely many experiments: Part I—Theory," Dept. of Statistics, Stanford Univ., Tech. Rep. 55, 1960.
- [15] M. Naghshvar and T. Javidi, "Active sequential hypothesis testing," *Ann. Stat.*, vol. 41, no. 6, pp. 2703–2738, 2013.
- [16] —, "Sequentiality and adaptivity gains in active hypothesis testing," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 5, pp. 768–782, Oct. 2013.

- [17] S. Nitinawarat, G. K. Atia, and V. V. Veeravalli, "Controlled sensing for multihypothesis testing," *IEEE Trans. Autom. Control*, vol. 58, no. 10, pp. 2451–2464, Oct. 2013.
- [18] D. Wei and A. O. Hero, "Multistage adaptive estimation of sparse signals," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 5, pp. 783–796, Oct. 2013.
- [19] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 3rd ed. Nashua, NH: Athena Scientific, 2005, vol. 1.
- [20] —, *Nonlinear Programming*. Belmont, MA: Athena Scientific, 1999.